

Rock Modeling and Matching for Autonomous Long-Range Mars Rover Localization

.....

.....

Rongxing Li and Kaichang Di

*Mapping and GIS Laboratory, CEEGS
The Ohio State University
Columbus, Ohio 43210
e-mail: li.282@osu.edu, di.2@osu.edu*

Andrew B. Howard and Larry H. Matthies

*Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California 91109
e-mail: abhoward@robotics.jpl.nasa.gov,
lhm@robotics.jpl.nasa.gov*

Jue Wang and Sanchit Agarwal

*Mapping and GIS Laboratory, CEEGS
The Ohio State University
Columbus, Ohio 43210
e-mail: wang.813@osu.edu, agarwal.59@osu.edu*

Received 9 June 2006; accepted 22 December 2006

This paper introduces the concept and design of a new integrated approach to long-range autonomous Mars rover localization based on the incremental bundle adjustment and visual odometry technologies that have been individually experimented with during the 2003 Mars Exploration Rover mission. The design result indicates that a rover would have a varying performance in traversing from 7.5 to 118 m within one traverse leg under various scenarios including camera systems and traverse geometry, while maintaining the onboard rover localization accuracy at 1%. To implement the proposed integrated approach, the key is to develop autonomous cross-site tie point selection algorithms for automatic generation of a sufficient number of high quality tie points to link all the images and to form the image network. New methods of rock extraction, rock modeling, and rock matching from multiple rover sites are developed to automate cross-site tie point selection. Rocks are extracted from three-dimensional ground points generated by stereo image matching, and then modeled using analytical surfaces such as hemispheroid, semiellipsoid, cone, and tetrahedron. Rocks extracted and modeled from two rover sites are matched by a combination of rock model matching and rock distribution pattern matching. The matched rocks are used as cross-site tie points for a subsequent bundle adjustment. The presented results show that the proposed cross-site tie point selection approach functions successfully for medium-range (up to 26 m) traverse legs. © 2007 Wiley Periodicals, Inc.

1. INTRODUCTION

In planetary surface robotic operations, such as the current Mars Exploration Rover (MER) mission and the planned 2009 Mars Science Laboratory (MSL) mission, it is critical to obtain the location and attitude information of a rover with a high degree of accuracy both for safe rover navigation and for achieving science and engineering goals (Arvidson et al., 2004; Li et al., 2004; Li, Squyres et al., 2005). For future explorations, the rover would be required to traverse more than the 9 km that is achieved by the Opportunity rover currently. There is always the desire to travel a longer distance within each command cycle so that the rover can explore a landing site as far as possible within its lifetime. Furthermore, the terrain to be explored may be so difficult that autonomous navigation would require highly detailed information of the environment in real time and at a high accuracy. In addition, the targets investigated may be very far from each other and the observations may be related spatially and/or temporally. These present a great challenge to rover localization and mapping technologies.

Mobile robot localization and navigation in an indoor or outdoor environment have been researched by Jarvis (1993), Borenstein, Everett, Feng & Wehe (1997) and Spero (2004). The positioning methodologies can be categorized as dead-reckoning methods (odometry and inertial navigation) and reference-based methods that includes Global Positioning Systems (GPS), landmark based navigation, and model (map) matching techniques. As demonstrated in the 2005 DARPA Grand Challenge, autonomous land vehicle navigation and localization achieved a successful level in an outdoor desert environment. Five vehicles completed the course of 132 miles between 6 h 53 min and 12 h 51 min (Buehler, 2006). However, these navigation and localization technologies may not be directly applicable to the Martian environment because of unavailability of GPS on Mars, payload and power limitation, and other Mars relevant constraints.

Planetary rover localization research has been carried out at the Jet Propulsion Laboratory (JPL) using several advanced methods including position and heading estimation by remote viewing of a colored cylindrical target (Volpe, Litwin & Matthies, 1995), maximum likelihood matching of range maps (Olson & Matthies, 1998; Olson, 2000), and visual odometry (VO) algorithms (Matthies, 1989; Olson,

Matthies, Shoppers & Maimone, 2000, 2003). Schenker, Huntsberger, Pirjanian, Baumgartner & Tunstel (2003) reviewed JPLs recent research on planetary rover developments that may support future Mars exploration and sample return missions. Many field tests were conducted by JPL on Earth to validate Mars rover technologies including navigation and localization (Baumgartner, Aghazarian & Trebi-Ollennu, 2001) and to rehearse planetary mission operations (Tunstel et al., 2002). The Robotics Institute at Carnegie Mellon University (CMU) designed and developed various robotic systems and vehicles for industrial and military applications. In recent years, CMU performed field experiments for robotic investigation of life in the Atacama Desert of Chile that is an analog to Mars (Wettergreen, Cabrol, Teza et al., 2005; Wettergreen, Cabrol, Baskaran et al., 2005). Rovers Hyperion and Zoë, respectively, were tested in the experiments, and long-range autonomous traverses (e.g., more than 1 km within a single command cycle) were achieved. The rover localization accuracy was 3%–5% of distance traveled based on dead-reckoning method that integrated wheel encoders, roll and pitch inclinometers and yaw gyro (Wettergreen, Cabrol, Teza et al., 2005). Centre National d'Etudes Spatiales is also developing Mars rover autonomous navigation technology based on Inertial Measurement Unit (IMU), odometry and stereo vision (Mauratte, 2003). Li et al. (2002, 2004) developed a bundle adjustment (BA) method for long range Mars rover localization using descent and rover images.

In the Mars Pathfinder (MPF) mission, the rover Sojourner was localized using dead-reckoning sensors. Daily location updates based on lander images were sent to the rover from Earth (Matthies et al., 1995). The overall localization error was found to be about 10% of the distance from the lander within an area of about 10 m × 10 m. This localization accuracy, while sufficient for the MPF mission, needed an improvement for longer range rover navigation.

In the MER mission, wheel odometry, sun finding technique using rover images, and IMU are used to estimate rover positions and attitudes with a designed accuracy of 10%. The onboard VO technique has been used to correct errors that are often caused by wheel slippage when the rovers travel on steep slopes or across loose soils (Maimone, Johnson, Cheng, Willson & Matthies, 2004). The BA method for rover localization and topographic mapping has been utilized on the ground (Earth) to improve the localization accuracy at the Gusev Crater and Meridiani

Planum landing sites (Li, Squyres et al., 2005; Li, Archinal et al., 2005; Di et al., 2005). It is demonstrated that the combined VO and BA methods are capable of correcting position errors caused by wheel slippage, azimuthal angle drift and other navigation errors as large as the 21% experienced within Eagle Crater at the Meridiani Planum landing site and the error of 10.5% in the Husband Hill area at the Gusev Crater landing site. Routinely produced orthoimages, digital terrain models (DTMs), and rover traverse maps have been provided to MER mission scientists and engineers through a web-based geographic information system that has greatly supported tactical and strategic mission operations (Li, Archinal et al., 2005).

This paper introduces the concept of a new approach to autonomous long-range Mars rover localization based on the integrated BA and VO method. Rover traverse design and theoretical analysis of achievable localization accuracy are discussed. Since the Spirit rover has successfully achieved an accuracy of 0.5% over a distance of 6 km using an onboard VO and the Earth version BA, automation of BA plays an important role for implementation of the proposed integrated BA and VO method for autonomous long-range localization. Tie point selection, particularly cross-site tie point selection is a key component for autonomous BA. We present a new approach to and the recent results of an automatic cross-site tie point selection method that includes rock extraction, rock modeling, and rock matching.

2. CONCEPT OF AN INTEGRATED APPROACH TO LONG-RANGE ROVER LOCALIZATION

2.1. Design of the Integrated Approach

This long-range rover localization technology is based on estimation of rover positions and attitudes at times of imaging through a three-dimensional (3D) image network consisting of overhead (orbital or descent) and rover imagery. The image network is created by linking the images using tie points which are commonly ground features, such as remote landmarks (e.g., hill peaks) and rocks, in the overlapping regions of the images. An incremental bundle adjustment of the image network provides highly accurate 3D positions of the tie points and rover position and attitude information through the reconstructed image pointing parameters (Li, Di et al., 2005; Li, Di,

Agarwal, Wang & Matthies, 2006). The proposed new autonomous rover localization method aims at three significant improvements over the Earth based rover localization technology used in the MER mission: (1) automation of the cross-site tie point selection process and development of an onboard version of the incremental BA software, (2) integration of the onboard incremental BA with VO to achieve long-range autonomous rover localization, and (3) development of an Earth-based BA software system that integrates both orbital and ground images of the entire image network to achieve the optimal localization accuracy and to update the rover position and attitude every few command cycles.

Within one command cycle which is usually one sol, the onboard incremental rover localization system is designed to perform real-time localization autonomously using rover images. VO and BA methods will be integrated with the expectation of achieving high efficiency and full automation. As illustrated in Figure 1, the rover travels from waypoint W_0 to W_2 within a sol, with a stop at midpoint W_1 . Usually, a full (360°) or partial panorama of Pancam (Panoramic Camera) or Navcam (Navigation Camera) images is taken at the start and endpoints. Images in the driving directions (forward and backward) are often acquired at the midpoint. BA is performed at the waypoints and the midpoint through the image network including images acquired before W_0 and at the incremental portion (W_0 , W_1 , and W_2). Further, VO is carried out between the waypoints using more densely collected sequential Navcam images. VO is used to correct accumulated errors in odometer-based measurements mainly caused by slippage and to provide improved positions and attitudes between BA-based updates. The VO results are then incorporated into BA in two ways: (a) VO-estimated incremental camera position and attitude information is used as the initial input to the incremental bundle adjustment and (b) the VO-tracked image features are passed to and employed in BA. Specifically, the VO-tracked features are directly used as tie points to link the sequential images together, while the cross-site tie point selection method presented in this paper is applied to link the images taken at waypoints and/or midpoints and also applied to link the first and last stereo pairs of the VO images to the closest waypoint/midpoint images. The onboard version of the incremental BA starts with the uplinked rover position and attitude of the last command cycle at a site




Figure 1. Configuration of the onboard image network. BA performed at the waypoints and the midpoint through the image network, while VO is carried out between the waypoints using sequential Navcam images.

location with a full or partial panorama (W_0), which can be computed on Earth at a higher accuracy. BA estimates the optimal rover positions and attitudes at locations within the current command cycle whenever the “incremental” panoramic and/or mid-point survey images are taken. The same process repeats in each command cycle.

The Earth-based BA is designed to be performed in two modules. First, on each sol an incremental bundle adjustment is carried out using rover images down linked during the previous three sols. Remote landmarks such as mountain peaks and craters extracted from orbital images are included to enhance long-range localization capabilities. The calculated rover location and attitude of the last location are uplinked each sol to provide the starting information of the following sol. Second, every three sols or longer, an integrated bundle adjustment of the entire image network consisting of all available rover and orbital images will provide the best rover location information. This Earth-based BA overcomes localization difficulties caused by wheel slippage as well as long-term accumulative errors such as IMU drift. Since the pointing information of each image in the network is improved as a byproduct of the BA process, high quality mapping products such as DTMs,

landmark maps, rover paths, and seamless orthoimages can then be generated to support navigation and other mission operations.

2.2. Rover Traverse Design and Accuracy Analysis

BA is a method that employs the above introduced image network along the long-range traverse for precision rover localization. Therefore, the configuration of the rover traverse directly affects the attainable localization accuracy, as well as the feasibility and effectiveness of the cross-site tie point selection method that is used to build the image network. To study the potential accuracy of rover localization under different traverse configurations, we performed a rover traverse design and accuracy analysis based on a rover traverse simulation study (Di, Li, Matthies & Olson, 2002). A number of variations including traverse leg length (distance between adjacent sites), number of tie points, and convergence angle (see definition later) were investigated. Table I lists the parameters of the cameras assumed in the analysis, including MER Navcam, MER Pancam, and a camera similar to the planned MSL Navcam. The range measurement error is proven to be proportional to the square of the range and the azimuth measurement error that can be set as the angular

Table I. Camera parameters used in localization accuracy analysis.

	MER Navcam	MSL-like Navcam	MER Pancam
Stereo base	20 cm	30 cm	30 cm
Focal length	14.67 mm	14.67 mm	43 mm

error caused by 1 pixel shift (Moffit & Mikhail, 1980; Di et al., 2002; Li & Di, 2005). The photogrammetric error propagation method is used to analyze the relationships between the desired rover localization accuracy (BA) and the various network configurations. Visual odometry data is not simulated nor used in this accuracy analysis. Since VO would provide better initial orientation parameters for images taken between waypoints, the overall accuracy of rover localization is ultimately determined by bundle adjustment of the image network. Thus, the estimated accuracy in this theoretical analysis should represent the attainable accuracy of the integrated rover localization approach.

The analysis was first performed for the case of a traverse with only two sites (Figure 2). The traverse leg is the distance between two sites (Site 0


**Figure 2.** Tie point distribution and convergence angle of a traverse with two sites. Cross-site tie points are evenly distributed in the middle of the overlapping area. Different convergence angles correspond to different relative spacing of the tie points.**Table II.** Optimal traverse parameters (convergence angle and leg length) estimated under various scenarios for maintaining one percent localization accuracy.

Image network configuration				
No. of tie points	Convergence angle and leg length	MER Navcam	MSL-like Navcam	MER Pancam
2	Angle (°)	71	71	71
	Length (m)	7.5	11.2	33
4	Angle (°)	88–86	90–86	86
	Length (m)	19	28.6	85
6	Angle (°)	86–84	88–84	84–83
	Length (m)	22	32.4	100
9	Angle (°)	79–64	86–66	64–62
	Length (m)	26	38.8	118

and Site 1 in Figure 2). We assume that the cross-site tie points which connect forward looking and backward looking images of the two sites are evenly distributed in the middle of the overlapping area. Figure 2 shows an example of nine tie points and the convergence angles which are defined as the angle at either site between the beginning and end sights covering the set of tie points. For each configuration, we fix the number of tie points and camera parameters, and change the convergence angle and the leg length to find the optimal convergence angle and traverse leg length for maintaining a desired localization accuracy (e.g., 1% for onboard localization). The results are illustrated in Table II. Specifically, the localization accuracy is given in a relative way, a ratio between the position error at the end of the traverse (Site 1) and the traverse length (distance between Site 0 and Site 1). For example, if the traverse length is 30 m and the position error at the end of the traverse is 30 cm, the (relative) localization accuracy is 1%. In Table II, the optimal convergence angle is the angle that will give the smallest localization error for different leg lengths with a fixed number of tie points and camera setting; it may be slightly different for different leg lengths (see the angle range in the table). Given a fixed number of tie points and camera setting, the optimal leg length is the maximum length of the traverse segment that meets the 1% error limit with an optimal convergence angle.

Based on the results of the two-site analysis in Table II, the following trends can be observed:

1. A desired rover localization accuracy can be achieved by different network configurations with various factors such as traverse leg length, convergence angle, number of tie points, and camera system parameters (stereo base, focal length, etc.).
2. Given other network parameters, increasing the number of tie points allows the rover to traverse a longer distance (leg length).
3. Given the number of tie points and other network parameters, a camera system with a longer stereo base and a longer focal length (Table I) can also support a longer traverse leg.
4. For example, with nine well-distributed tie points in the middle of two sites, a rover can limit its localization error within 1% and at the same time increase its traverse leg length from 26 m (using MER Navcam) to 39 m (using MSL-like Navcam), or even to 118 m (using MER Pancam).

To extend the two-site case to the entire traverse with many sites, we applied the camera parameters in Table I and the analysis results of the two-site traverse segment in Table II to propagate the errors through the traverse (Li & Di, 2005). Under the constraints of nine tie points and the localization accuracy of 1% for each traverse leg (one segment), we computed the required number of sites, traverse leg length, and the expected localization accuracy at the end of a 5 km traverse as presented in Table III.

Suppose for one traverse segment with a leg length of D_1 , the absolute localization error (standard deviation) is m_1 , the relative error is then m_1/D_1 . According to the error propagation rule, the absolute localization error of an n segment traverse will be $\sqrt{nm_1}$ if all the legs have the same length of D_1 (overall traverse length nD_1) and the same error of m_1 . As a result, the relative localization error of the entire traverse will be $\sqrt{nm_1}/(nD_1) = m_1/(\sqrt{n}D_1)$, which is $1/\sqrt{n}$ of the relative error of one traverse segment. That is why we can observe that the relative localization accuracy of a 5 km traverse is better than those of one leg in Table III.

From Table III, we can observe that if the relative localization accuracy of each traverse leg is better than 1%, an overall relative localization accuracy of

Table III. Traverse design of and error analysis for a 5 km traverse.

Camera system	MER Navcam	MSL-like Navcam	MER Pancam
No. of sites	194	130	44
Traverse leg length	26	38.8	118
No. of tie points within one leg	9	9	9
Localization accuracy of one leg (%)	1	1	1
Localization accuracy of the 5 km traverse (%)	0.072	0.088	0.15

0.1% to 0.2% would be achievable for a 5 km traverse. The derivation and computation process of Table III is lengthy and cannot be fully described in this paper because of the space limit. The technical details are discussed in Li & Di (2005). Furthermore, this result is verified to be close to the performance of the Spirit rover at Gusev crater on Mars where the rover achieved an accuracy of 0.5% over a distance of 6 km during its initial operation period without the above assumptions met completely. We expect that post mission processing will result in an accuracy that will be more comparable to that in Table III although some traverse leg lengths are beyond those specified in Table III.

In the development of the automatic cross-site tie point selection method, we considered the above theoretical analysis by picking the test data according to the optimal leg length and selecting evenly distributed tie points in the middle of the overlapping area. Our goal is to automate the cross-site tie point selection process while maintaining one percent accuracy for each traverse segment.

2.3. Visual Odometry

As illustrated in Figure 1, VO is critical for rover position and attitude updates between waypoints where BA uses the VO result as its initial input into the model and provides the improved rover location and attitude information. In addition, VO also contributes to the achievement of the desired accuracy

(e.g., 1%) between two adjacent sites (one traverse leg, Figure 2) in the above traverse accuracy analysis, so that the overall traverse accuracy (e.g., 0.1%–0.2%) can be obtained. In principle, the VO algorithm estimates the rover motion by tracking visual features between consecutive stereo pairs in both two-dimensional (2D) pixel coordinates and 3D ground coordinates (Matthies, 1989; Olson et al., 2000, 2003). It has been successfully used in the MER mission for precision instrument placement, estimation of local rover position and attitude information by overcoming wheel slippage, and supply of the initial input to BA for global localization (Maimone et al., 2004; Li, Squyres et al., 2005). Further improvement of the VO algorithm is being carried out to increase reliability, accuracy and driving speed. We are investigating an inertial-visual-odometry algorithm that fuses data from three sensors: a stereo camera system, an IMU, and wheel encoders.

3. AUTOMATIC CROSS-SITE TIE POINT SELECTION FOR AUTONOMOUS INCREMENTAL BUNDLE ADJUSTMENT

Implementation of the proposed integrated approach to long-range rover localization depends on automation of VO and BA performed onboard the rover. VO has been successfully employed onboard the rovers in the MER mission. BA has been effectively applied in the mission as a ground software system because most of the cross-site tie points are selected by human operators. Therefore, the key is to develop autonomous cross-site tie point selection algorithms for automatic generation of a sufficient number of high quality tie points to link all the images and to form the image network. The selection of the tie points within one full panorama at one site has been automated and routinely performed in the MER rover localization (Xu, 2004; Di et al., 2005). A great challenge is the automatic selection of cross-site tie points where objects (e.g., rocks) used as tie points look significantly different from different views, especially from forward and backward views. We developed a new approach to automatic selection of rocks as cross-site tie points through rock modeling and rock matching as described in Figure 3.

Rocks are one of the major features/objects found in rover images of the Martian terrain. In order to link rover images taken at adjacent sites and thus build the image network, it is desirable to extract, model

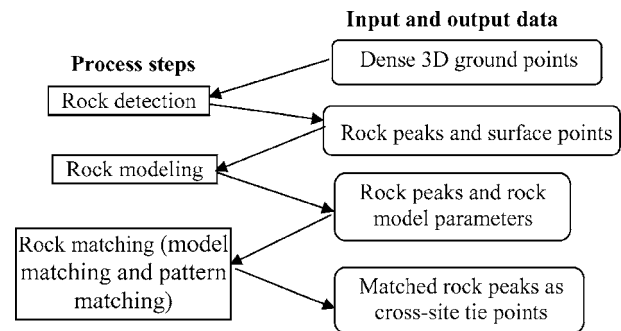


Figure 3. Diagram of automatic cross-site tie point selection.

and match rocks from multiview ground images. Object detection and modeling from multiple view-points has been addressed in the computer vision community. A complete survey of object detection and geometric modeling techniques is outside the scope of our paper. We only briefly review several most relevant papers below.

Gor, Castano, Manduchi, Anderson & Mjolsness (2001) developed a rock detection method that uses image intensity information to detect small rocks and uses range information to detect large rocks from Mars rover images. In detection of large rocks, a least-squares plane (ground) is fitted using the range data, and the above-ground height of a range point is calculated as the distance between that point to the plane. Large rocks are then extracted through a height segmentation process. Subsequently, the shape of the extracted rocks is modeled by seven metrics: eccentricity, ellipse error, 2D sphericity, 2D angularity, ellipsoid error, 3D sphericity, and 3D angularity (Castano et al., 2002; Fox, Castano & Anderson, 2002). The rock detection and shape modeling methods are included in JPL's Onboard Autonomous Science Investigation System and have been tested for use in autonomous rock shape analysis by the Mars rover, the aim of which is to enable the rover to direct its activities to regions of higher geologic significance and to prioritize the downlink of data, thus maximizing the scientific return of the mission (Castano et al., 2005).

CMU researchers developed a rock detection method based on segmentation, detection, and classification using texture, color, shape, shading, and stereo data from the Zoë rover (Thompson, Niekum, Smith & Wettergreen, 2005). A multiple-view detec-

tion method was also developed by the same group to fuse rocks extracted from multiple views based on expectation maximization clustering, which does not require an explicit correspondence search of rocks (Thompson & Wettergreen, 2005). The aim of this multiple-view object detection is to yield a better understanding of the scene than any one view alone. The rover images of different views used in the experiment were acquired from relatively close positions (2 m apart).

In our approach, the rocks are detected from dense ground points generated from stereo images acquired at each site (Figure 3). As a result, the rock detection process provides a rock peak and a set of rock surface points for each rock, which are further used to fit a number of analytical surface models such as cone or spheroid. Each rock is then represented by a surface model that best fits the extracted rock surface points. After the rocks are modeled, matching of rocks from adjacent sites are carried out by a comparison of individual rock models as well as the two global rock (peak) distribution patterns from the two sites. Finally, the matched rock peaks are utilized to link rover images and to build the image network.

Our rock detection and modeling methods are different from the existing methods mentioned above, reflecting the different purposes of the methods. The purposes of the existing methods are to detect as many rocks as possible and to model the shapes of the rocks for autonomous geological analysis. The objectives of our methods are to extract and match a sufficient number of rocks from different views at two adjacent rover sites (which are far apart) for autonomous rover localization. It is not necessary to extract small rocks as long as the number of big rocks is sufficient. As for rock modeling, from our experience of manual cross-site tie point selection, the size of a rock is usually more significant than the shape of the rock for multi-view matching. Thus, the purpose of our rock modeling (surface fitting) method is to reliably estimate the size (i.e., height, radius, etc.) of a rock.

3.1. Rock Extraction

In order to detect and describe rocks, detailed topographic information of the area between two sites needs to be reconstructed. To a large degree, general terrain features can be characterized by “interest” points which could be rocks peaks, sharp corners, ridge points, and others. They are extracted from




Figure 4. TIN-controlled 3D dense image matching. The vertices of the triangles are the matched and verified interest points on the left image of a stereo pair. The triangles were created and used to predict the approximate location of the homologous point of a dense grid point. The white dots are matched dense grid points.

stereo images at each site using a Förstner interest operator. The interest points are then matched by using crosscorrelation. They are further processed by verification of parallax consistency and outlier elimination. These interest points generally cover most rock peaks that are sufficient for BA (Xu, 2004; Di et al., 2005).

The matched interest points are practically too sparse to model the rocks, particularly in the middle or far range from the rover. To obtain potential rock surface points, a triangulated irregular network (TIN) controlled dense image matching is performed to produce a terrain model (grid) with a grid-cell size of 3×3 or 5×5 pixels. First, a TIN is generated from the matched interest points on the left image (Figure 4); the parallax is then calculated at each node (vertex). Second, for each grid point on the left image, its homologous image point on the right image is estimated using linear interpolation of parallax from the vertices of the triangle that covers the grid point. The actual homologous point in the right image is refined by crosscorrelation within a small search range (e.g., three pixels) from the estimated position. Since the dense matching is well controlled within a very small window controlled by the reliably matched interest point, this two-level matching strategy improves the overall reliability and speed of




Figure 5. Rock peak (white dot) and surface points (gray dots) extracted in the iterative process for one rock.

image matching. After dense matching, dense 3D ground points are calculated through space intersection of homologous image points.

Figure 4 shows an example of dense matching using MER Navcam images taken by the Spirit rover. It can be observed that there are a sufficient number of 3D ground points (both interest points and grid points) on the rock surfaces to facilitate rock extraction and modeling in the subsequent steps. We can also see that there are some grid points missing, where the cross-correlation coefficient is less than a predefined threshold (e.g., 0.75).

Rock peaks are extracted from the reconstructed 3D ground points using the following criteria: (1) they are the local maxima within a window of, for example, 50×50 cm; (2) the maximum elevation difference within the window is greater than a threshold (e.g., 10 cm); and (3) there are a sufficient number of ground points (at least three) within the window.

Extraction of Rock Surface Points

In addition to the rock peak, rock surface points are needed to describe the shape of a rock and also to fit an analytical rock model. Initiating from a rock peak, we look for ground points that are on the same rock in the vicinity of the peak by an automatic searching procedure.

1. A 3D plane is estimated using the terrain points within an area of $70 \text{ cm} \times 70 \text{ cm}$ from the rock peak. In principle, the candidate rock surface points are those above the plane. On the other hand, the ground points are on or

very close to the plane. The rock height H is calculated as the perpendicular distance from the peak to the fitted plane.

2. Surface points are searched among the points above the fitted plane using a dynamic search range. The search range is proportional to the rock height H : kH , where k varies from 0.3 to 1.7 based on a ground truth experiment in which manual measurements of rocks at the Spirit site were made and the coefficient k was calculated. Initially the search from the rock peak is made in a range with the low limit radius of $0.3H$. Each time a point is found, the distance from this point to the closest neighboring point is calculated. The maximum of these distances, MaxD_Neighbor , is determined.
3. An extended search is performed from each point found in the last step and within the range of MaxD_Neighbor . Therefore, the overall search radius from the peak is increased by MaxD_Neighbor . This step repeats for all points found in the last step until the overall search radius reaches the upper limit of $1.7H$, or no new points in the neighborhood can be added.

Figure 5 shows the rock peak and rock surface points of one rock extracted in the iterative process. The white dot is the rock peak, while the gray dots are the extracted surface points. Additional examples of extracted rock peaks and surface points of other rocks are shown in Figure 6.

The above method was successfully applied to extract peaks and surface points of various types of




Figure 6. Other examples of extracted rock peaks (white dots) and rock surface points (gray dots) of six different rocks.

rocks that are about 15 m from the rover, including large rocks (e.g., 0.5 m high). The algorithm met difficulties when dealing with a rock complex where a number of rocks stand closely together and severe occlusions block some rock peaks and surface points. However, as long as a sufficient number of rocks can be extracted and matched between two sites, the incremental BA can be achieved.

3.2. Rock Modeling

Compared to the above discrete information of the rock peak and surface points extracted from images taken at a certain view point(s), an analytical model, such as a cone, provides more information under the assumption that the peak and surface points can be used to infer the part of the rock surface which is not visible to the camera if a symmetric model is used. Furthermore, a rock model with its parameters is more efficient for comparison with other rocks modeled in the same way.

Based on the extracted rock peak and surface points, we can model a rock by using a 3D analytical surface, such as hemispheroid, semiellipsoid, cone, or tetrahedron. The equations for a hemispheroid, a semiellipsoid, and a cone are

$$\frac{x^2}{r^2} + \frac{y^2}{r^2} + \frac{z^2}{h^2} = 1, \quad (1)$$

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} + \frac{z^2}{h^2} = 1, \quad (2)$$

and

$$x^2 + y^2 = (z - h)^2 \frac{r^2}{h^2}, \quad (3)$$

respectively. The parameter r is the radius of the hemispheroid or radius of the bottom circle of the cone; a and b are semimajor and semiminor axes of the semiellipsoid; and h is the height in all three equations. No analytical equation exists for the tetrahedron model. Therefore we use three parameters to represent a tetrahedron: height h , radius of the enclosing circle of the bottom triangle, and orientation angle ϕ of the bottom triangle.

The model parameters of each individual rock are estimated by a least-squares fitting using the surface points on the rock. Since the equations of these analytical models are nonlinear with respect to the model parameters, linearized equations are derived using Taylor's theorem and are used for least squares computation, where the model parameters are calculated iteratively.

As an example, we give the mathematical details of the hemispheroid model fitting. Rewriting Eq. (1) as $F(r, h) = x^2/r^2 + y^2/r^2 + z^2/h^2 - 1 = 0$, the linearized hemispheroid equation becomes

$$F = F_0 + \frac{\partial F}{\partial r} \Delta r + \frac{\partial F}{\partial h} \Delta h = 0, \quad (4)$$

where r and h are unknowns; $\partial F / \partial r = -2(x^2 + y^2)/r^3$ and $\partial F / \partial h = -2(z^2/h^3)$ are partial derivatives; F_0 is the value of $F(r, h)$ calculated using the initial values of r and h . For example, the value of rock height H obtained in the previous step of rock surface point detection can now be used as the initial value of h in Eq. (4). Then, the observation equation for least squares fitting is

$$\begin{bmatrix} \frac{\partial F}{\partial r} & \frac{\partial F}{\partial h} \end{bmatrix} \begin{bmatrix} \Delta r \\ \Delta h \end{bmatrix} = -F_0. \quad (5)$$

Each surface point contributes to one observation equation of Eq. (5). Putting the observation equations for all n surface points ($i=1, 2, \dots, n$) of a rock together, a matrix form of the observation equations can be expressed as

$$AX = L \quad (6)$$

$$\text{where } A = \begin{pmatrix} -\frac{2(x_1^2 + y_1^2)}{r^3} & -\frac{2z_1^2}{h^3} \\ -\frac{2(x_2^2 + y_2^2)}{r^3} & -\frac{2z_2^2}{h^3} \\ \vdots & \vdots \\ -\frac{2(x_n^2 + y_n^2)}{r^3} & -\frac{2z_n^2}{h^3} \end{pmatrix}_{n \times 2},$$

$$L = \begin{pmatrix} 1 - \frac{x_1^2}{r^2} + \frac{y_1^2}{r^2} + \frac{z_1^2}{h^2} \\ 1 - \frac{x_2^2}{r^2} + \frac{y_2^2}{r^2} + \frac{z_2^2}{h^2} \\ \vdots \\ 1 - \frac{x_n^2}{r^2} + \frac{y_n^2}{r^2} + \frac{z_n^2}{h^2} \end{pmatrix}_{n \times 1} \quad \text{and} \quad X = \begin{bmatrix} \Delta r \\ \Delta h \end{bmatrix}.$$

The least squares solution of Eq. (6) is

$$X = (A^T A)^{-1} A^T L.$$

In j th iteration, the model parameters are updated by

$$\begin{bmatrix} r \\ h \end{bmatrix}_j = \begin{bmatrix} r \\ h \end{bmatrix}_{j-1} + \begin{bmatrix} \Delta r \\ \Delta h \end{bmatrix}_j. \quad (7)$$

Points that are far (e.g., 10 cm) from the estimated surface are eliminated during the iterative process. If the number of the remaining surface points is less than 5 after elimination, the computation is terminated and the rock is considered not to fit the particular model. If none of the models fits a rock, the rock will be discarded.

To evaluate the fitting accuracy of each model for a particular rock, the root mean square (rms) error of the fitted surface model is calculated

$$rms = \sqrt{\frac{(z_i - z_{i,model})^2}{n}}, \quad (8)$$

where z_i is the height value of the i th surface point, $z_{i,model}$ is the corresponding height value calculated using the fitted model.

Figure 7 shows nine rocks of different sizes and shapes. Each rock is modeled by using the four analytical models. Table IV gives the estimated model parameters of the four models for each rock, as well as the associated rms errors of the fitting models. The highlighted (italic) rms error of each rock is the minimum among the four models. The corresponding model is considered the best model for that rock. For example, Rock 1 is best represented by a semiellipsoid, while Rock 2 is best represented by a tetrahedron. In addition to statistical evaluation using rms error (as shown in Table IV), the fitted model can be backprojected onto the image to check the




Figure 7. Nine sample rocks with marked peaks (white dots) used in rock modeling. The best-fit tetrahedron model of the ninth rock can be seen overlaid on its image.

quality of model fitting. As an example, the tetrahedron model (best fit) of the ninth rock can be seen overlaid on its image in Figure 7.

To verify the rock modeling results, we compared the modeled parameters with the ground truth (manual measurements from stereo images) of 79 rocks in the area between two adjacent sites that are 26 m apart at the Spirit landing site. For each rock four metrics were used, including height, radius, surface area, and volume, which were calculated from the ground truth measurements and the estimated model. For the semiellipsoid model, for instance, the average of a and b is considered as the radius in the comparison. In ground truth measurements, we measured the rock height directly; we measured the rock width and used half of the width as the radius. The ground truth surface area and volume are calculated from the measured height and radius. On average, of the 79 rocks, the relative difference between a modeled parameter and the corresponding ground truth measurement is 25.1% in height, 43.7% in radius, 57.1% in surface area, and 103.4% in volume, respectively. It is obvious that among the four metrics, height is the most reliable one and thus the most comparable parameter. The very high difference in volume suggests that it should not be used for comparison. This verification result is important for designing the following algorithm of rock matching with various models.

3.3. Rock Matching

Rock matching is used to find corresponding rocks in the two sets of extracted rocks from two sites. There are various difficulties in this process: (1) rocks visible from one site may not be visible or

Table IV. Estimated model parameters and rms errors (ϕ in radians, all others in cm) of rocks in Figure 7.

Rock. ID	Hemispheroid model			Cone model			Semiellipsoid model				Tetrahedron model			
	a	h	rms	r	h	rms	a	b	h	rms	a	h	ϕ	rms
1	12.2	15.5	4.5	17.0	16.7	3.93	11.1	14.6	15.6	3.92	14.0	16.7	1.2	4.1
2	20.8	26.8	5.1	23.5	28.5	4.9	20.1	23.2	27.4	5.2	22.5	28.5	1.0	4.4
3	26.3	23.1	5.0	30.6	24.6	4.8	21.0	32.0	24.0	5.4	32.6	24.6	1.2	3.8
4	9.3	13.4	5.4	5.7	13.8	5.4	17.0	5.2	12.4	4.8	8.7	13.8	1.3	4.1
5	9.1	9.6	2.0	19.0	9.4	1.7	10.6	6.8	9.7	1.52	22.0	9.4	1.2	1.51
6	40.9	35.3	5.1	31.6	40.5	5.8	38.5	59.4	34.1	5.1	34.6	40.5	0.1	3.4
7	22.6	20.0	4.7	25.8	22.3	5.2	14.7	27.3	20.3	4.6	22.8	22.3	0.8	3.7
8	28.7	24.2	5.5	30.9	25.4	4.1	25.4	31.3	24.4	5.2	34.9	25.4	0.8	4.5
9	13.5	13.2	4.0	16.2	14.6	3.4	15.1	13.0	13.2	3.9	22.2	14.6	0.8	2.9

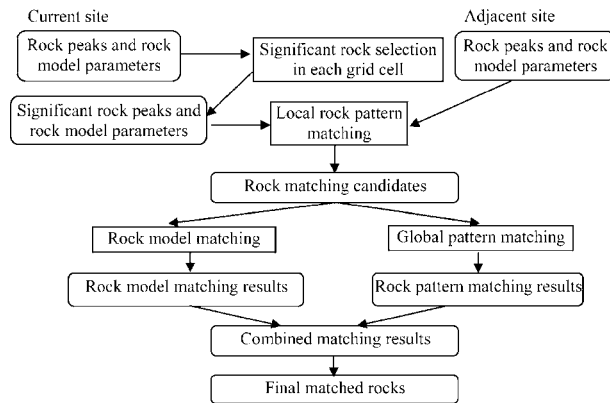

identifiable from the other site because of occlusions and significant distance differences from the sites; (2) the same rock extracted from different views may not have the same surface portion imaged; and (3) information extracted from smaller rocks may not be sufficiently reliable for matching because of the limited stereo range capability.

The rock matching technique we developed considers both global rock distribution patterns by rock pattern matching and individual rock similarities by rock model matching. Figure 8 shows a detailed diagram of the entire rock matching process.

The two sites are called the current site and adjacent site in Figures 8 and 9. Rock extraction and modeling are carried out separately from the images of the two sites. Rock peaks and rock model parameters are estimated for each rock. The question is, for

one rock extracted at the current site, which one extracted at the adjacent site is its correspondent. We define a 4×4 grid in the overlapping area between the two sites (Figure 9). In implementation, the overlapping area is determined by an intersection of two $60 \text{ m} \times 60 \text{ m}$ boxes surrounding the two site centers. Within each grid cell, we select a limited number of significant rocks (e.g., up to 3), which are usually the highest rocks in the grid cell. Note that some of the grid cells may not have any significant rocks either because there are no rocks in a grid cell or none of the extracted rocks in the cell can be modeled appropriately. Only the selected significant rocks at the current site are used to find their corresponding rocks at the adjacent site.

For each significant rock extracted at the current site, a set of extracted rocks at the adjacent site are individually compared by the rock model matching technique. The model matching for one rock of the current site and another of the adjacent site uses an objective function

**Figure 8.** Rock matching technique using both model matching and rock distribution pattern matching.**Figure 9.** A grid is defined in the middle between the two sites for extraction of cross-site tie points.

$$Z = c_1 f_1 + c_2 f_2 + c_3 f_3, \quad (9)$$

where f_1 , f_2 , and f_3 are the relative differences (in percentage) of height, radius, and surface area between the two rocks calculated from the two rock models. The coefficients c_1 , c_2 , and c_3 are the relevant weights, which are set to be $1/2$, $1/3$, and $1/6$ based on the result of rock model verification described above. The height is given the largest weight since it is the most comparable parameter. The rock of the adjacent site with the minimum value of Z in Eq. (9) is considered a match.

On the other hand, rock pattern matching compares the two geometric distributions of the rock peaks at the current and adjacent sites. In principle, a rigid transformation including three rotations and a 3D translation can be used to depict the relationship between two rock distribution patterns derived from the two sites. Based on extensive experiments using Spirit rover data, it was found that the rotational differences are insignificant. This may be because the sun-positioning technique was periodically used for absolute rover heading detection during mission operations. Consequently, a 3D translation is employed in rock pattern matching. Furthermore, this makes the pattern matching process computationally more efficient. Assume that we have two sets of extracted rock peaks extracted from the current site and adjacent site with their positions as P_i ($i=1,2,\dots,I$) and Q_j ($j=1,2,\dots,J$), respectively, in the 3D ground coordinate system. For rock pattern matching, a subset of significant peaks P_k ($k=1,2,\dots,K$) are selected from the points P_i ($i=1,2,\dots,I$) and are used to find matches with Q_j ($j=1,2,\dots,J$). The pattern matching algorithm works as follows:

1. For a significant peak P_k from the current site, we define a searching area (e.g., within a distance of about 20% of the traverse leg length) where we find a subset Q_m ($m=1,2,\dots,M$) out of Q_j from the adjacent site. For each candidate peak Q_m at the adjacent site, calculate a translation vector $V_{k,m}=P_k-Q_m$.
2. Apply the translation vector to all the rock peaks Q_j at the adjacent site $Q'_j=Q_j+V_{k,m}$.
3. We check the consistency between P_k and Q_m through a pattern comparison using the translation vector $\text{Count}_{k,m}=\text{Card}(\|P_i-Q'_j\| \leq \varepsilon)$, where ε is a predefined tolerance. $\|\cdot\|$ denotes the norm of the difference vector. Card

function represents the number of the pairs of rock peaks that can be potentially matched.

4. Repeating steps (1)–(3) for all the points in Q_m , we have all the counts for the point P_k .
5. For each significant peak P_k at the current site, two best candidate rocks Q_{m1} and Q_{m2} in Q_m at the adjacent site are found, which have the highest and second highest counts with their corresponding translation vectors $V_{k,m1}$ and $V_{k,m2}$. We refer to the above process as local pattern matching. We repeat the local pattern matching for all significant rocks in P_k .
6. For global pattern matching, we check the consistency of all translation vectors of the significant rocks P_k ($k=1,2,\dots,K$). We determine the vector V_{median} that represents the median among all the translation vectors. Then, for each significant rock P_k , find its final match by checking the difference between the individual translation vector $V_{k,m1}$ and the global translation vector V_{median} . Consequently, Q_{m1} is considered as a match of P_k if $\|V_{k,m1}-V_{\text{median}}\| \leq \|V_{k,m2}-V_{\text{median}}\|$; otherwise P_k matches with Q_{m2} . This global matching is performed for all significant rocks.

Studies using various data sets of Spirit rover found that a significant number of the correct matches correspond to the second highest counts in the step 3 above. That is why in the local pattern matching we select two candidate rocks from the adjacent site for each significant rock at the current site. It should also be noted that the local pattern matching improves the efficiency and robustness of the rock model matching technique by limiting to two rock candidates and reducing the number of possible mismatches in model matching.

In both rock model matching and rock pattern matching, there exist multiple-match cases. That is, two or more rocks from the current site match the same rock from the adjacent site. This indicates mismatches although one match among these multiple matches might be a correct match. The rocks with multiple matches are automatically eliminated in the matching processes.

It was also found that the results of rock model matching and rock pattern matching are not very different. In the end the final matching result is the combination of the outputs of the two methods. Only the rocks that pass both matching methods are considered to be the matched rock pairs.




Figure 10. Extracted rocks from two sites of Spirit rover. The triangles are the rock extracted from Site 9600, while the crosses are those extracted from Site 9700. The rover location of Sites 9600 and 9700 are marked as dots.

Figure 10 shows an example of automatically extracted rocks for cross-site tie point selection. Navcam images were taken at Site 9600 and Site 9700 by the Spirit rover. The two sites are 26 m apart. Thirty-seven rocks are extracted from Site 9600 (triangles) and 34 rocks from Site 9700 (crosses). Nine significant rocks were automatically selected from the rocks of Site 9600 in a grid defined between the two sites. Table V is a summary of the experimental results of rock matching, including rock model matching, rock pattern matching, and the combined method for comparison purposes. Figure 11 shows the final matched rock pairs (combined method), which are labeled with the same identification numbers for both sites.

In Table V, for each particular rock, the rock model matching (RMM) algorithm and the rock pat-




Figure 11. Automatically matched rocks as cross-sites tie points that are labeled with the same identification numbers. The triangles are the rocks extracted from Site 9600, while the crosses are those extracted from Site 9700. The rover locations of Sites 9600 and 9700 are marked with dots.

tern matching (RPM) algorithm generate one of two outputs: a unique match, or elimination of the rock because of multiple matches. The final matching result is generated by the combination of RMM and RPM. For each rock, the final output is either a unique match (if RMM and RPM matching results are the same) or no match (if RMM and RPM matching results are different or the rock is eliminated because of multiple matches). In order to verify whether the matched rocks are correct, the nine rocks in Table V were also manually matched by an operator and the manual matching results are used as ground truth. By comparing the matching results with ground truth, the unique match is validated as a correct match or mismatch in Table V. From Figure

Table V. Summary of rock matching results for Sites 9600 and 9700.

Method/Rock ID	1	2	3	4	5	6	7	8	9
Model matching (8 correct matches)	– ^d	√ ^a	√	√	√	× ^b	√	√	√
Pattern matching (7 correct matches)	–	√	√	√	√	×	√	√	√
Combined Method (7 correct matches)	O ^c	√	√	√	√	O	√	√	√

^a√ Correct match.
^b× Mismatch.
^cO No match.
^d– Eliminated because of multiple matches.

Table VI. Summary of rock matching results for Sites 11493 and 11422.

Method/Rock ID	1	2	3	4	5	6	7	8	9	10	11	12	13
Model matching (8 correct matches)	√ ^a	× ^b	√	×	√	√	√	√	√	×	– ^d	√	×
Pattern matching (9 correct matches)	×	√	√	√	√	√	√	√	√	×	×	√	×
Combined Method (7 correct matches)	O ^c	O	√	O	√	√	√	√	√	O	O	√	O

^a√ Corrected match.^b× Mismatch.^cO No match.^d– Eliminated because of multiple matches.

11 and Table V, it is demonstrated that rock model matching (after local rock pattern matching) found eight matched rock pairs correctly, while one was eliminated because of multiple matches. Pattern matching in Table V was performed by local pattern matching followed by global pattern matching. It produced seven matched rock pairs. But there were one mismatched rock and a rock eliminated. In the final rock matching result, seven rock pairs were found correctly and two were unmatched. The seven correctly matched tie points are sufficient for BA of the two sites.

Table VI gives another example of automatic cross-site tie point selection using Navcam images taken at Site 11493 and Site 11422 by the Spirit rover. The two sites are 18 m apart and the terrain presents relatively steep slopes. In the combined result, 7 out of 13 rocks were matched by the software and they are all correct matches. No mismatches were created in the combined result. The distribution of the matched rocks is also good for BA because they are well distributed in cells of the defined grid. From both Table V and VI we can observe that rock model matching and pattern matching generated slightly different results. Their combinations eliminated the mismatches and made the results reliable (no mismatches).

4. CONCLUSIONS

This paper introduces the concept and design of a new integrated approach to long-range autonomous Mars rover localization. In the design, we looked at changing various rover/camera/environment parameters and estimated how far the rover could

traverse while maintaining 1% onboard localization accuracy. The design result indicates that a rover would have a varying performance in traversing from 7.5 to 118 m to within one traverse leg under various scenarios to meet the accuracy requirement. To implement the proposed integrated localization approach, automatic cross-site tie point selection is a key process and is developed based on rock extraction, rock modeling, and rock matching using images acquired at multiple rover sites. In rock matching, the rock model matching and rock pattern matching methods can be employed complementally to eliminate potential mismatches and to enhance reliability of the matching results. The presented cross-site tie point selection results show that the proposed approach functions successfully for medium-range (up to 26 m) traverse legs. The presented traverse design and the achieved cross-site tie point selection results will be used in the future to integrate BA and VO methods for autonomous onboard long-range rover localization.

ACKNOWLEDGMENTS

Part of this work is being performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Funding of this research by the Mars Technology Program of NASA/JPL is acknowledged. The authors appreciate constructive comments from the reviewers. Thanks to Irene Tesfai and Leslie Smith of The Ohio State University for proofreading several versions of the manuscript.

REFERENCES

- Arvidson, R.E., Anderson, R.C., Bartlett, P., Bell, J.F. III, Blaney, D., Christensen, P., et al. (2004). Initial localization and physical properties experiments conducted by Spirit at Gusev Crater. *Science*, 305(5685), 821–824.
- Baumgartner, E., Aghazarian, H., & Trebi-Ollennu, A. (2001, October). Rover localization results for the FIDO rover. Paper presented at the SPIE Photonics East Conference, Boston, MA.
- Borenstein, J., Everett, H.R., Feng, L., & Wehe, D. (1997). Mobile robot positioning: Sensors and techniques. *Journal of Robotic Systems*, 14(4), 231–249.
- Buehler, M. (2006). Summary of DGC 2005 results. *Journal of Field Robotics*, 23(8), 465–466.
- Castano, R., Anderson, R.C., Fox, J., Dohm, J., Haldemann, A.F.C., & Fink, W. (2002, March). Automating shape analysis of rocks on Mars. Paper presented at the 33rd Lunar and Planetary Science Conference, Abstract No. 2000, Lunar and Planetary Institute, Houston, TX.
- Castano, R., Judd, M., Estlin, T., Anderson, R.C., Gaines, D., Castano, A., et al. (2005, March). Current result from a rover science data analysis system. In *Proceedings of the 2005 IEEE Aerospace Conference, Big Sky, MT*.
- Di, K., Li, R., Matthies, L.H., & Olson, C.F. (2002, April). A study on optimal design of image traverse networks for Mars rover localization. In *Proceedings of the ACSM-ASPRS 2002 Annual Conference, Washington, D.C.*
- Di, K., Xu, F., Wang, J., Niu, X., Serafy, C., Zhou, F., et al. (2005, March). Surface imagery based mapping and rover localization for the 2003 Mars Exploration Rover mission. In *Proceedings of the ASPRS 2005 Annual Conference, Baltimore, MD*.
- Fox, F., Castano, R., & Anderson, R.C. (2002, March). On-board autonomous rock shape analysis for Mars rovers. In *Proceeding of the 2002 IEEE Aerospace Conference, Big Sky, MT*.
- Gor, V., Castano, R., Manduchi, R., Anderson, R.C., & Mjolsness, E. (2001, August). Autonomous rock detection for Mars terrain. Paper presented at AIAA Space 2001—Conference and Exposition, Albuquerque, NM.
- Jarvis, R. (1993, July). A selective survey of localisation methodology for autonomous mobile robot navigation. In *Proceeding of the Australian Robot Association's International Conference on Robots for Competitive Industry, Brisbane, Australia* (pp. 310–321).
- Li, R., Archinal, B.A., Arvidson, R.E., Bell, J., Christensen, P., Crumpler, L., et al. (2005). Rover localization and topographic mapping at the landing site of Gusev Crater Mars. *Journal of Geophysical Research—Planets*, 111, E02S06.
- Li, R., & Di, K. (2005, February). Rover traverse design and error analysis. Technical report submitted to JPL Mars Technology Program. Columbus, OH: Mapping and GIS Laboratory, CEEGS, The Ohio State University.
- Li, R., Di, K., Agarwal, S., Wang, J., & Matthies, L. (2006, March). Object modeling and matching from multi-view ground images for automated Mars rover localization. In *Proceedings of the IEEE Aerospace Conference, Big Sky, MT*. DOI 10.1002/rob
- Li, R., Di, K., Agarwal, S., Wang, J., Matthies, L., Howard, A., & Willson, R. (2005, September). Incremental bundle adjustment techniques using networked overhead and ground imagery for long-range autonomous Mars rover localization. In *Proceedings of the ISAIRAS 2005 Conference, Munich, Germany*.
- Li, R., Di, K., Matthies, L.H., Arvidson, R.E., Folkner, W.M., & Archinal, B.A. (2004). Rover localization and landing site mapping technology for 2003 Mars Exploration Rover mission. *Photogrammetric Engineering and Remote Sensing*, 70(1), 77–90.
- Li, R., Ma, F., Xu, F., Matthies, L.H., Olson, C.F., & Arvidson, R.E. (2002). Localization of Mars rovers using descent and surface-based image data. *Journal of Geophysical Research—Planets*, 107(E11), FIDO 4.1–4.8.
- Li, R., Squyres, S.W., Arvidson, R.E., Archinal, B.A., Bell, J., Cheng, Y., et al. (2005). Initial results of rover localization and topographic mapping for the 2003 Mars Exploration Rover mission. *Photogrammetric Engineering and Remote Sensing*, 71(10), 1129–1142.
- Maimone, M., Johnson, A., Cheng, Y., Willson, R., & Matthies, L. (2004, June). Autonomous navigation results from the Mars Exploration Rover (MER) mission. In *Proceedings of the 9th International Symposium on Experimental Robotics (ISER), Singapore*.
- Matthies, L.H. (1989). Dynamic stereo vision. Unpublished Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA.
- Matthies, L., Gat, E., Harrison, R., Wilcox, B., Volpe, R., & Litwin, T. (1995). Mars microrover navigation: Performance evaluation and enhancement. *Autonomous Robots*, 2(4), 291–311.
- Mauratte, E. (2003). Mars rover autonomous navigation. *Autonomous Robots*, 14(2–3), 199–208.
- Moffitt, F.H., & Mikhail, E.M. (1980). *Photogrammetry* (3rd ed.). New York: Harper and Row.
- Olson, C. (2000). Probabilistic self-localization for mobile robots. *IEEE Transactions on Robotics and Automation*, 16(1), 55–66.
- Olson, C., & Matthies, L. (1998, May). Maximum likelihood rover localization by matching range maps. In *Proceedings of the IEEE International Conference on Robotics and Automation, Lueven, Belgium*.
- Olson, C.F., Matthies, L.H., Shoppers, M., & Maimone, M. (2000, June). Robust stereo egomotion for long distance navigation. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Hilton Head Island, SC*.
- Olson, C.F., Matthies, L.H., Shoppers, M., & Maimone, M. (2003). Rover navigation using stereo egomotion. *Robotics and Autonomous Systems*, 43(4), 215–229.
- Schenker, P.S., Huntsberger, T.L., Pirjanian, P., Baumgartner, E.T., & Tunstel, E. (2003). Planetary rover developments supporting Mars exploration, sample return and future human-robotic colonization. *Autonomous Robots*, 14(2–3), 103–126.
- Spero, D.J. (2004, November). A review of outdoor robotics research (Tech. Rep. MECSE-17–2004). Melbourne, Australia: Department of Electrical and Computer

- Systems Engineering, Monash University.
- Thompson, D.R., Niekum, S., Smith, T., & Wettergreen, D. (2005, March). Automatic detection and classification of geological features of interest. In Proceedings of the IEEE Aerospace Conference, Big Sky, MT.
- Thompson, D.R., & Wettergreen, D. (2005, August). Multiple-object detection in natural scenes with multiple-view expectation maximization clustering. Paper presented at the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2005), Edmonton, Canada.
- Tunstel, E., Huntsberger, T., Aghazarian, H., Backes, P., Baumgartner, E., Cheng, Y., et al. (2002, June). FIDO rover field trials as rehearsal for the 2003 Mars Exploration Rover mission. Paper presented at the 9th International Symposium on Robotics and Applications, 5th World Automation Congress, Orlando, FL.
- Volpe, R., Litwin, T., & Matthies, L.H. (1995, August). Mobile robot localization by remote viewing of a colored cylinder. In Proceedings of the International Conference on Robots and Systems (IROS), Pittsburgh, PA.
- Wettergreen, D., Cabrol, N., Baskaran, V., Calderon, F., Heys, S., Jonak, D., et al. (2005, September). Second experiments in the robotic investigation of life in the Atacama Desert of Chile. In Proceedings ISAIRAS 2005 Conference, Munich, Germany.
- Wettergreen, D., Cabrol, N., Teza, J., Tompkins, P., Urmson, C., Verma, V., et al. (2005, April). First experiments in the robotic investigation of life in the Atacama Desert of Chile. Paper presented at the IEEE International Conference on Robotics and Automation, Barcelona, Spain.
- Xu, F. (2004). Mapping and localization for extraterrestrial robotic explorations. Unpublished Ph.D. dissertation, The Ohio State University, Columbus, OH.
- Xu, F., Di, K., Wang, J., & Li, R. (2005, March). Techniques of registration and mapping in Mars Exploration Rovers (MER) 2003 mission. In Proceedings of the ASPRS 2005 Annual Conference, Baltimore, MD.